

16719ROUS01U

DATA TRAFFIC POLICER

Field of the Invention

5 The present invention relates to data traffic policers and is particularly concerned with token or leaky bucket policers.

Background of the Invention

10 Data traffic may include several forwarding classes (emission priorities) that require different treatment on access to the data network. Data may also be assigned various levels of discard priority, sometimes designated by colour, for example green, yellow and red for high, medium and low priority, respectively. Both of these priorities need to be taken into account by data network entities involved in traffic management. Such entities include schedulers, data shapers, and policers. While schedulers and data shapers effect changes in the makeup of the data stream, policers monitor the data stream and identify violations of the constraints imposed on the data network, for example throughput. Because of this difference, policers provide an opportunity to ensure fairness with respect to network resource utilization.

20 Known policers have tried to use a serial or a hierarchical configuration to provide for both emission priority and discard priority. However in such configurations information can be lost between policers, thereby adversely affecting the policer's fairness. Hence, it would be desirable to provide a policer that maintains fairness.

Summary of the Invention

25 An object of the present invention is to provide an improved data traffic policer.

In accordance with an aspect of the present invention there is provided a data traffic policer comprising a classifier for separating a packet stream in accordance

16719ROUS01U

with class, a first bucket for a first traffic class representing a first transmission rate and a first burst capacity and a second bucket for a second traffic class representing a second transmission rate and a second burst capacity, the second bucket being nested within the first bucket thereby being subordinate to the rate and capacity of the first bucket, with the rate of the second bucket being disabled when a fill condition exists in the first bucket.

In accordance with another aspect of the present invention there is provided a method of data policing comprising the steps separating a packet stream in accordance with class, representing a first traffic class as a first transmission rate and a first burst capacity, and representing a second traffic class as a second transmission rate and a second burst capacity being subordinate to the rate and capacity of the first traffic class, with the rate of the second traffic class being disabled when a fill condition exists for the first traffic class.

Brief Description of the Drawings

The present invention will be further understood from the following detailed description with reference to the drawings in which:

Fig. 1 illustrates a data traffic policer in accordance with an embodiment of the present invention; and

Fig. 2 illustrates in a block diagram the data traffic policer of Fig. 1 with a different fill state.

Detailed Description of the Preferred Embodiment

Referring to Fig. 1, there is illustrated in a block diagram a data traffic policer in accordance with an embodiment of the present invention. The data traffic policer includes a first leaky bucket 12 having a rate R and a plurality of second leaky buckets 14, 16, 18 and 20 for receiving respective portions of data traffic from an input 22 via a corresponding plurality of classes 24, 26, 28, 30 and 32. The first leaky

16719ROUS01U

bucket 12 is fed by a committed traffic class 24, labeled green (G) for convenience. Second leaky buckets 14, 16, 18, 20 are fed by respective forwarding classes $FC_1, FC_2, FC_3, \dots, FC_n$. The green traffic directly fills the first leaky bucket, as illustrated with a green traffic fill level F_c . The first leaky bucket 12 has a fill limit B_c 34, visually represented by the upper edge of the bucket.

Similarly, each of the second leaky buckets 14, 16, 18, 20 has a respective limit B_{si} , however they also have a collective limit B_e (for excess), such that: $B_e = \sum B_{si}$, and $B_e < B_c$. Each of the second leaky buckets may also have separate limits for each of the discard classes (colours), for example B_{ri} for red and B_{yi} for yellow. Thus, for second bucket 14, the bucket limit for red B_{r1} may be the bucket upper edge 36, while bucket limit for yellow B_{y1} may be a lower limit 44. Similarly, for the second bucket 16, the bucket limit for red B_{r2} may be the bucket upper edge 38, while bucket limit for yellow B_{y2} may be a lower limit 46. Similarly, for the second bucket 18, the bucket limit for red B_{r3} may be the bucket upper edge 40, while bucket limit for yellow B_{y3} may be a lower limit 48. Similarly, for the second bucket 20, the bucket limit for red B_{rn} may be the bucket upper edge 42, while bucket limit for yellow B_{yn} may be a lower limit 50. Each of the second buckets also has assigned a weight W_i that is used to determine the rate at which they leak, once the committed traffic has been satisfied.

Operation of the data traffic policer is described with reference to Figs. 1 and 2. In operation, traffic received on data input 22 is split according to class. Committed traffic 24 is applied directly to the first leaky bucket 12. As long as F_c is non-zero, the rate R is completely consumed by the committed traffic. Thus, as shown in Fig. 1, the fill levels of the second buckets 14-20 does not go down and likely increases with $F_c > 0$. Once the committed traffic has been satisfied, as illustrated in Fig. 2, the second leaky buckets 14 - 20 can begin to leak at their respectively weighted rates. Consequently, the respective fill levels $F_1, F_2, F_3, \dots, F_n$, begin to lower, however, as soon as any committed (green) traffic 24 arrives, it is applied directly to the first leaky bucket 12, F_c becomes non-zero and the second buckets 14 - 20 are prevented from further emptying.

The above description is a view of how the data traffic policer in accordance with embodiments of the present invention can be considered conceptually. Implementation of this view may be in the form of an algorithm applied to the data network to effect the data traffic policer. The following tables, Table A and Table B, provide the static configuration parameters and dynamic parameters needed in an implementation of the policer algorithm.

TABLE A Static Configuration Parameters

Symbol	Definition
R	Rate of the aggregate policer
B_c	Bucket limit for committed (green) traffic
B_e	Bucket limit for non-committed traffic $B_e < B_c$
W_i	Weight for each FC = 1,2...n
B_{yi}	Bucket limit for yellow traffic FC = 1,2...n
B_{ri}	Bucket limit for red traffic FC = 1,2...n

Note: B_{yi} could be configured as a percentage of B_{ri} , thereby only requiring a single parameter for all FCs.

TABLE B Dynamic Parameters

F_c	Fill level of committed (green) traffic
F_i	Fill level of excess (non-green) traffic of FCi
F	Fill level of aggregate policer ($F_c + \text{sum}(F_i)$)

Input:

Per-packet: Colour (Green, Yellow or Red), FC (1 .. N), Packet Size L in bytes

T_c = Current Time

Static Configuration:

R = Rate of the aggregate policer

16719ROUS01U

Here Rate may be Bytes per millisecond (therefore 8Mbps translates to
 $R=1000$ Bytes per millisecond)

B_c = Bucket limit for green (committed) traffic

B_e = Bucket limit for non-green (excess) traffic ($B_e < B_c$)

5 W_i = Weight of each FC, $i = 1..N$

B_{yi} = Bucket limit for yellow traffic of FC, $i = 1..N$

B_{ri} = Bucket limit for red traffic of FC, $i = 1..N$

10 Implementation note: B_{yi} could be implemented as a percentage of B_{ri} ,
 thereby requiring only a single parameter for all FCs. A further simplification would
 be to set B_{ri} the same as B_{yi} .

Dynamic Parameters:

F_c = Fill level of green (committed) traffic

15 F_i = Fill level of non-green (excess) traffic of FC_i

F = Fill level of aggregate policer (same as F_c + all F_i)

T = Last time packet was received (for example stored in milliseconds). An
 actual implementation may use microseconds for greater accuracy.

All fill parameters are initialized to 0.

20

Algorithm:

Leakage of Buckets:

25 The aggregate fill level is decremented at the rate of R . That is: decremented
 by $D=R*(T_c-T)$.

D is divided between the F_c and F_i 's as follows:

If $F_c \geq D$, $F_c = (F_c - D)$ and F_i 's are unchanged. In other words, as long
 as F_c is non-zero the rate R is applied to F_c .

30 If $F_c < D$, $F_c = 0$, and the residue $(D - F_c)$ is divided between the F_i 's
 according to their weight.

16719ROUS01U

Given the above condition of applying R first to F_c , what this means is that for some time T_0 , the rate R was applied to F_c until $F_c = 0$, then from T_0 until time T_c , R was applied to the remaining buckets F_i . How this is done is implementation specific (for e.g., a bulk WRR could be used).

5

The following example is provided to demonstrate the leaky bucket algorithm.

Let $T = 0$, $F_c = 10 \text{ kB}$ and $R = 1000 \text{ B/ms}$

a) at $T_c = 1 \text{ ms}$ $F_c = F_c - D$

10

$$\begin{aligned} F_c &= 10 \text{ kB} - 1 \text{ kB} \\ &= 9 \text{ kB} \quad \text{yes } F_c > D \end{aligned}$$

b) at $T_c = 5 \text{ ms}$ $D = 1000 \times (5 - 0) = 5 \text{ kB}$

15

$$\begin{aligned} F_c &= F_c - D \\ F_c &= 10 \text{ kB} - 5 \text{ kB} \\ &= 5 \text{ kB} \quad \text{yes } F_c > D \end{aligned}$$

c) at $T_c = 10 \text{ ms}$ $D = 1000 \times (10 - 0) = 10 \text{ kB}$

20

$$\begin{aligned} F_c &= F_c - D \\ F_c &= 10 \text{ kB} - 10 \text{ kB} \\ &= 0 \text{ kB} \quad \text{yes } F_c = D \end{aligned}$$

d) at $T_c = 11 \text{ ms}$ $D = 1000 \times (11 - 0) = 11 \text{ kB}$ $F_c < D$

25

$$\begin{aligned} F_c &= 0 \\ \text{Residue} &= D - F_c = 11 \text{ kB} - 10 \text{ kB} \\ &= 1 \text{ kB} \quad \text{to be applied to bucket fill levels } F_i \text{ according to weights } W_i. \end{aligned}$$

30

Hence, the committed traffic uses the entire leaky bucket rate R as long as F_c is non-zero. However once the fill decrement D is greater than F_c , any residue $D - F_c$ is used to decrement the other buckets with fill levels F_i

16719ROUS01U

When a new packet arrives, T is updated with value of T_c (i.e: $T=T_c$) and the buckets are incremented in accordance with the following.

Incrimination of Buckets:

Green packet: If $F < B_c$, /* Packet is conforming */

5

{

$F = F + L$;

$F_c = F_c + L$;

}

/* Else packet is non-conforming */

10

Yellow packet, Class i : If $F < B_e$ AND $F_i < B_{yi}$ /* Packet is conforming */

{

$F = F + L$;

$F_i = F_i + L$;

15

}

/* Else packet is non-conforming */

Red packet, Class i : If $F < B_e$ AND $F_i < B_{ri}$ /* Packet is conforming */

{

20

$F = F + L$;

$F_i = F_i + L$;

}

/* Else packet is non-conforming */

/* End of HWP description */

25

If a packet is nonconforming, various options are available with regard to determining the behavior of the policer. For example, the policer may mark the packet (in which case it is counted in the bucket) for discard in the event of downstream congestion conditions. Alternatively, the policer may immediately discard the packet, in which case it would not contribute to the bucket count.

30

16719ROUS01U

While the embodiment of the present describe herein above uses leaky buckets, it will be appreciated by those of ordinary skill in the art that an alternative embodiment could be provided using token buckets.

5 Also in the figures a two-level nested bucket hierarchy has been presented for simplicity of the description, however it should be appreciated that a plurality of levels is also possible.

10 Numerous modification, variations and adaptations may be made to the particular embodiments of the invention described above without departing from the scope of the invention, which is defined in the claims.